


Guidance for Advertising Agents on AWS


This architecture diagram demonstrates how to automate complex advertising campaign planning workflow using Generative AI Agents on AWS. This slide shows steps 1-



- Business Users securely authenticate through **Amazon Cognito** user pools. The AngularJS UI application, hosted in **Amazon Simple Storage Service (Amazon S3)** and distributed via **Amazon CloudFront** for low-latency global access, validates these **Amazon Cognito** tokens to authorize user sessions.
- User asks for Agent help in complex topics like planning a premium product campaign or optimizing inventory yield in natural language using chat interface, calling out the agent's name using @<agent> prefix.
- The query is passed to an Intelligent Orchestrator Strands Agent deployed in **Amazon Bedrock AgentCore** Runtime. Based on the input one of the Orchestrator agents, with agent depending on the invocation prompt, Media Planner, Inventory Optimizer, or Ad Load Optimizer handles the query.
- On initialization, the Orchestrator agent uses an agent config file packaged in the Agent Core run time to identify other specialized agents it needs to collaborate with. The specialized agents are then configured using Strands agent as a tool framework. The Orchestrator offers dual interaction modes - users can engage with the orchestration agents for comprehensive response, or directly chat with specialized agents for quick, focused answers.
- The Orchestrator agent uses **Amazon Bedrock** Knowledge Base indexed with campaign intelligence, performance analytics, content safety, and viewer monetization data stored in **Amazon S3**.
- The orchestrator agents and specialized agents collaboratively work together in natural language conversation, each focusing on areas like audience strategy, inventory optimization, ad format selection, and campaign timing.

Guidance for Advertising Agents on AWS

This architecture diagram demonstrates how to automate complex advertising campaign planning workflow using Generative AI Agents on AWS. This slide shows steps 1-



- 7 The agents use visualization template files stored in **Amazon S3** to transform complex data to assist rendering of visual elements.
- 8 The inter-agent conversations and the final response is streamed back to the user in real time. The final response contains integrated campaign strategies with interactive visualizations, including audience segments insights and revenue projections.
- 9 The conversation history between the user and agents is stored in **Amazon Bedrock AgentCore** Memory.
- 10 **Amazon CloudWatch** collects performance metrics and operational data from **Amazon Bedrock AgentCore** observability features, including session information and conversation traces.
- 11 **AWS Identity and Access Management (AWS IAM)** implements fine-grained permissions for service-to-service communication and **Amazon Bedrock AgentCore** access following a least-privilege access model.